"AI doomer" is a label for someone who is concerned about AI-caused extinction (or comparable AI catastrophes) — often used pejoratively, by people who take a dim view of such concerns.

It's worth distinguishing a few kinds of views that might cause someone to be called an "AI doomer":
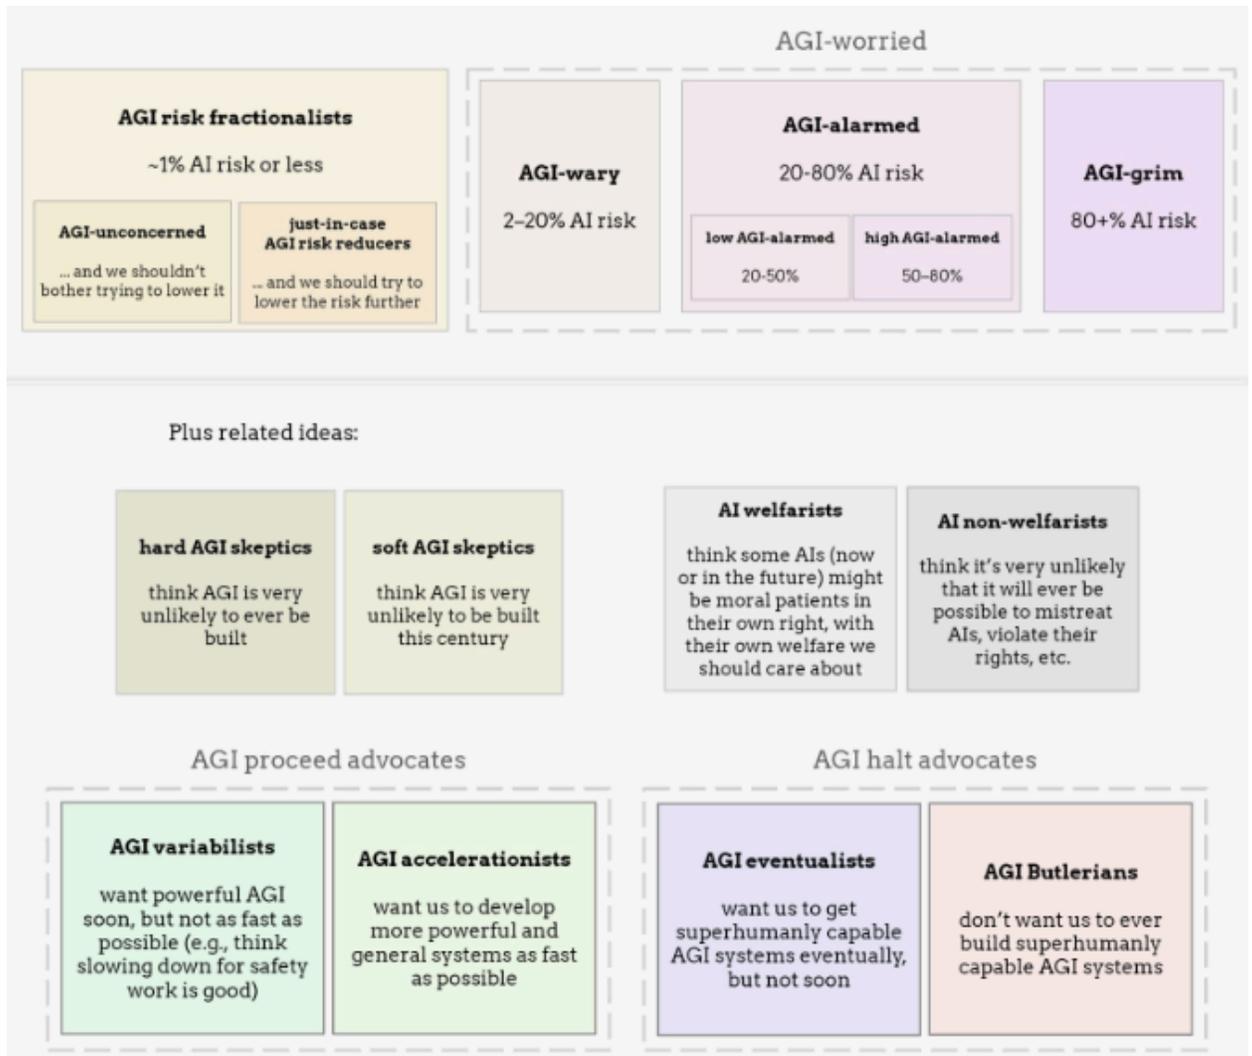
1. "AI-caused extinction is **plausible** enough to pay attention to and to make tradeoffs to prevent." This view is held by many AI researchers, engineers, policy analysts, safety researchers, and heads of AI labs, as illustrated by the [CAIS statement on AI risk](#).
2. "AI-caused extinction is highly **probable**." People with this view generally think that although we will probably fail to avoid extinction, it's worth trying. This includes [Eliezer Yudkowsky](#) and [others](#) at [MIRI](#).
3. "AI-caused extinction is **inevitable** and not even worth trying to prevent." This sense of "doomer" is often used in other contexts, like [climate change](#), but it's uncommon for people to be fatalistic about AI risk in this way.

Outside the context of existential risk, the term is occasionally applied to ideas as varied as "AI will doom artists to unemployment" and even "AI is doomed to stay at a low capability level." And outside the context of AI, "doomer" can refer to anyone who is pessimistic about technological progress (like a [Luddite](#)), or the future in general. This results in further confusion because many "AI doomers" are optimistic about progress, but carve out an exception for AI.

## Fixing the Ambiguity

While there are no alternatives to "AI doomer" that have gained traction, one attempt at better-defined terms comes from [Rob Bensinger](#):

- "AGI risk fractionalists": [p(doom)](#) < 2%
- "AGI-wary": p(doom) around 2-20%
- "AGI-alarmed": p(doom) around 20-80%
- "AGI-grim": p(doom) > 80%

AGI-worried

**AGI risk fractionalists**

~1% AI risk or less

| **AGI-unconcerned** | **just-in-case AGI risk reducers** |
|---|---|
| ... and we shouldn't bother trying to lower it | ... and we should try to lower the risk further |

**AGI-wary**

2–20% AI risk

**AGI-alarmed**

20-80% AI risk

| **low AGI-alarmed** | **high AGI-alarmed** |
|---|---|
| 20-50% | 50–80% |

**AGI-grim**

80+% AI risk

Plus related ideas:

**hard AGI skeptics**

think AGI is very unlikely to ever be built

**soft AGI skeptics**

think AGI is very unlikely to be built this century

**AI welfarists**

think some AIs (now or in the future) might be moral patients in their own right, with their own welfare we should care about

**AI non-welfarists**

think it's very unlikely that it will ever be possible to mistreat AIs, violate their rights, etc.

AGI proceed advocates

AGI halt advocates

**AGI variabilists**

want powerful AGI soon, but not as fast as possible (e.g., think slowing down for safety work is good)

**AGI accelerationists**

want us to develop more powerful and general systems as fast as possible

**AGI eventualists**

want us to get superhumanly capable AGI systems eventually, but not soon

**AGI Butlerians**

don't want us to ever build superhumanly capable AGI systems

Bensinger also suggests terms like "AGI variabilists" and "AGI eventualists" for different preferred policies about whether and when to build AGI. The policy question is partly separate from that of how likely AI is to cause extinction, but the term "doomer" tends to conflate them.

Further reading:
- [Some reasons not to say "Doomer"](#) by Ruby
- Ajeya Cotra, Shane Legg and others [discussing the relevance of the term](#)

## Alternative phrasings
- What is "AI doomerism"?

## Related

- ▤ What is "p(doom)"?
- ▤ How likely is extinction from superintelligent AI?

**Scratchpad**

See also  ▤ How likely is extinction from superintelligent AI?

alternative bullet draft (or maybe just some points to be integrated)

- it's an ambiguous and often pejorative term for people who worry about existential risk ("doom") from AI, sometimes also called ["safetyists", "x-risk advocates", …]
- in some other contexts, e.g. climate, "doomer" is used by people who stress that there's a problem but that it's worth solving, to refer to people who focus on it being catastrophic regardless of what we do; but this is mostly not how it's used in the AI context, so beware of confusion here
  - the word "doom" itself is somewhat ambiguous in that it can sometimes just mean "ruin" / "permanent bad event" / "extinction", but it can also imply inevitability; "p(doom)" [afaik] always means "probability of extinction" and not "probability that extinction is inevitable"
- in the AI x-risk context, "doomer" is often used to refer to anyone who thinks doom is plausible enough to be worth worrying a lot about, e.g. Hinton and Bengio get called "doomers" though their probability of x-risk is [look up] 10% or so
- an intermediate meaning is that a "doomer" is someone who thinks doom will probably happen but thinks it's not certain and worth trying to avoid even though we'll probably fail; Eliezer Yudkowsky and MIRI are the most prominent examples

general vague definition: "people sometimes use the word doomer to refer, mostly pejoratively, to those who believe AI poses an existential risk"

- "doomer" sense 1: someone who thinks it's plausible enough to pay attention to and make tradeoffs for (e.g. everyone who signed the CAIS letter)
- "doomer" sense 2: someone who thinks it's highly probable (e.g. MIRI)
- "doomer" sense 3: someone who thinks it's inevitable (not used much in AI risk discourse, but worth calling out as a possible misinterpretation because it's common in e.g. climate discourse
https://news.harvard.edu/gazette/story/2024/04/forget-doomers-warming-can-be-stopped-top-climate-scientist-says/
https://www.worldwildlife.org/stories/what-you-need-to-know-about-climate-doomers)